arXiv:1701.00299v1 [cs.LG] 2 Jan 2017

Dynamic Deep Neural Networks: Optimizing Accuracy-Efficiency Trade-offs by Selective Execution

Lanlan Liu llanlan@umich.edu
Jia Deng jiadeng@umich.edu
University of Michigan
2260 Hayward St, Ann Arbor, MI, 48105, USA

Abstract

We introduce Dynamic Deep Neural Networks (D^2NN), a new type of feed-forward deep neural network that allow selective execution. Given an input, only a subset of D^2NN neurons are executed, and the particular subset is determined by the D^2NN itself. By pruning unnecessary computation depending on input, D^2NNs provide a way to improve computational efficiency. To achieve dynamic selective execution, a D^2NN augments a regular feed-forward deep neural network (directed acyclic graph of differentiable modules) with one or more controller modules. Each controller module is a sub-network whose output is a decision that controls whether other modules can execute. A D^2NN is trained end to end. Both regular modules and controller modules in a D^2NN are learnable and are jointly trained to optimize both accuracy and efficiency. Such training is achieved by integrating backpropagation with reinforcement learning. With extensive experiments of various D^2NN architectures on image classification tasks, we demonstrate that D^2NNs are general and flexible, and can effectively optimize accuracy-efficiency trade-offs.

1. Introduction

This paper introduces Dynamic Deep Neural Networks (D^2NN), a new type of feed-forward deep neural network that allows selective execution. That is, given an input, only a subset of neurons are executed, and the particular subset is determined by the network itself and dependent on the particular input. In other words, the amount of computation and the sequence of computation are dynamic based on input. This is different from standard feedforward networks that always execute the same sequence of computation regardless of input.

A D^2NN is a regular feed-forward deep neural network (directed acyclic graph of differentiable modules) augmented with one or more control modules. A control module is a sub-network whose output is a decision that controls whether other modules can execute. Fig. 1 illustrates a simple D^2NN with one control module (Q) and two regular modules (N1 and N2), where the controller Q outputs a binary decision on whether module N2 executes. For certain inputs, the controller may decide that N2 is unnecessary and instead execute a dummy node to save on computation. As an example application, this D^2NN can be used for binary classification of images, where some images can be rapidly classified as negative after only a small amount of computation.

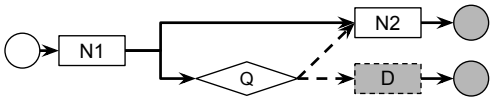

Figure 1. A simple D^2NN example. A control subnetwork Q decides whether to execute the subnetwork N2 or end the execution (by executing a dummy node D).

D^2NNs are motivated by the need for computational efficiency, in particular, by the need to deploy deep networks on mobile devices and data centers. Mobile devices are constrained by energy and power, limiting the amount of computation that can be executed in total and per unit time. Data centers need energy efficiency to scale to higher throughput and to save operating cost. D^2NNs provide a way to improve computational efficiency by selective execution, pruning unnecessary computation depending on input. In particular, D^2NNs make it possible to use a bigger network under a computation budget by executing only a subset of the neurons each time.

A D^2NN is trained end to end. That is, regular modules and control modules are jointly trained to optimize both accuracy and efficiency. We achieve such training by integrating backpropagation with reinforcement learning, necessitated by the non-differentiability of control modules.

Compared to prior work that optimizes computational efficiency in computer vision and machine learning (e.g. [23,

30, 2]), our work is distinctive in three aspects: (1) the decisions on selective execution are part of the network inference and are learned end to end together with the rest of the network, as opposed to hand-designed or separately learned [23, 30, 29, 2]; (2) D^{2}NNs allow more flexible network architectures and execution sequences including parallel paths, as opposed to single thread linear chains [11]; (3) our learning method directly optimizes accuracy metrics such as the F-score that does not decompose over individual examples. This is an issue not addressed in prior work. We will elaborate on these differences in Sec. 2 of this paper.

We perform extensive experiments to validate our D^{2}NN algorithms. We evaluate several different D^{2}NN architectures on image classification tasks including both binary and multiclass classification. Experiments demonstrate that D^{2}NNs are general, flexible, and can effectively improve computational efficiency.

Our main contribution is the D^{2}NN framework that allows a user to augment a regular feed-forward network with control modules to achieve dynamic selective execution. We show that D^{2}NNs allow a wide variety of topologies while sharing a unified training algorithm. Our D^{2}NN framework thus provides a new tool for designing and training computationally efficient neural network models.

2. Related work

Input-dependent execution has been widely used in computer vision, from cascaded detectors [32, 13, 27] to hierarchical classification [9, 5, 24]. The key difference of our work from prior work is that we *jointly* learn both visual features and control decisions *end to end*, whereas prior work either hand-designs features and control decisions (e.g. thresholding), or learns them separately.

In the context of deep networks, two lines of prior work have attempted to improve computational efficiency. One line of work tries to eliminate redundancy in data or computation in a way that is input independent. The strategies include approximating layers with simpler functions [12, 34, 21], using number representations of limited precision [7, 14, 16], and pruning network parameters [18, 22, 17]. The other line of work exploits the fact that not all inputs require the same amount of computation, and explores input-dependent execution of DNNs. Our work belongs to the second line of work, and we will contrast our work mainly with them.

Among methods leveraging input-dependent execution, some use pre-defined execution-control policies. For example, cascade methods [23, 30, 29] rely on manually-selected thresholds to control execution; Dynamic Capacity Network [2] uses a specially designed method to calculate a saliency map that is used to control execution. Our D^{2}NNs, instead, are fully learn-able; the execution-control policies of D^{2}NNs do not require manual design and are learned together with the rest of the network.

Our work is closely related to "conditional computation" methods [6, 4], which activate part of a network depending on input with the goal of improving efficiency. In particular, they learn policies to encourage sparse activations for stochastic neurons. Our work differs from these methods in several ways. First, our control policies are learned to directly optimize arbitrary user-defined global performance metrics including ones that can be defined only on a set of input examples (e.g. F score), whereas conditional computation methods have only learned policies that optimize for sparsity of local neural activations. In addition, D^{2}NNs allow more flexible topologies and parametrization. For example, in prior work [4], every neuron (or every block of neurons) is coupled with a control policy parametrized as the sigmoid of an affine transformation of the activations from the layer below, whereas in our case the policy can take arbitrary input and have arbitrary parametrization. Finally, D^{2}NNs augment regular feedforward networks and are completely deterministic, whereas the conditional computation networks use stochastic neurons and stochastic policies.

Our work is also related to attention models [10, 25, 28, 15, 33, 3, 2], some of which can perform input-dependent computation by selecting certain parts of the input for processing. Note that attention models can be categorized as "hard" attention [25, 3, 2] versus "soft" attention [15, 28]. "Hard attention" models only process the "salient" parts and discard others (e.g. processing only a subset of image sub-windows); in contrast, "soft attention" models process all parts but up-weight the "salient" parts (e.g. processing all image sub-windows although the weight for each window is input-dependent). Thus only hard attention models perform input-dependent execution as D^{2}NNs do. However, hard attention models differ from D^{2}NNs because hard attention models have typically involved only one attention module whereas D^{2}NNs can have multiple attention (controller) modules. In other words, conventional hard attention models are "single-threaded" whereas D^{2}NN can be "multi-threaded" and thus more flexbile. In addition, prior work in hard attention models has not directly optimized for accuracy-efficiency trade-offs.

D^{2}NNs also bear some similarity to Deep Sequential Neural Networks (DSNN) [11]. A DSNN is a directed acyclic graph where each node is a regular neural network layer. The inference of a DSNN starts from the root node of the graph and follows the directed edges until a dead end. Whenever there are multiple outgoing edges at a node, one edge is picked stochastically using a learned policy — given an input, we walk along the graph visiting a sequence of nodes. The main difference between D^{2}NNs and DSNNs is that DSNNs are "single-threaded": the computation in DSNNs always consists of exactly one path in the graph,

whereas for D^{2}NNs it is possible to have multiple paths or even the entire graph activated. Another difference is that DSNNs are stochastic whereas D^{2}NNs are deterministic. Last but not the least, although DSNNs can in principle be used to optimize accuracy-efficiency trade-offs, no such empirical results have been reported.

3. Definition and Semantics of D^{2}NNs

In this section we precisely define a D^{2}NN and describe its semantics, i.e. how a D^{2}NN performs inference.

D^{2}NN definition A D^{2}NN is a directed acyclic graph consisting of nodes and directed edges. There are three types of nodes: input nodes, output nodes, and function nodes. An input or output node represents a real-value vector as an input or output of the network. A function node represents a (differentiable) function that maps a vector to another vector. There are two types of edges: data edges and control edges. A data edge represents a vector sent from one node to another, the same as in a conventional feedforward network. A control edge represents a control signal, a scalar, sent from one node to another. For a data edge originating from a function node, it can optionally have a user-defined "default value", representing the output that will still be sent even if the function node did not execute.

We place a few restrictions on the placement of data edges and control edges: (1) between any two nodes there can only be one edge, either a data edge or a control edge, but not both; (2) an input node has only outgoing data edges and an output node has one and only one incoming data edge; (3) the outgoing edges from a node are either all data edges or all control edges (i.e. cannot be a mix of data edges and control edges); (4) if a node has an incoming control edge, it cannot have an outgoing control edge. Note that these constraints are for simplicity of implemention, and do not in any way restrict the expressiveness of a D^{2}NN. For example, to achieve the effect of a node with a mix of outgoing data edges and control edges, we can just feed its data output to a new node with outgoing control edges and let the new node compute the identity function.

We call a function node a *control node* if its outgoing edges are control edges. We call a function node a *regular node* if its outgoing edges are data edges. Note that it is possible for a function node to take no data input and output a constant value if executed. We call such nodes "dummy" nodes. Hereafter we may also call function nodes "layers", "subnetwork", or "modules" and will use these terms interchangeably.

Fig. 2 illustrates a simple D^{2}NN with all kinds of nodes and edges.The function nodes are drawn as rectangles or diamonds depending on whether they are regular nodes or control nodes. The input and output nodes are drawn as circles with the output nodes shaded. Dummy nodes are shaded to indicate zero computation. Data edges are drawn as solid arrows and control edges as dashed arrows. A data edge with a user defined default value is drawn as a solid arrow decorated with a circle.

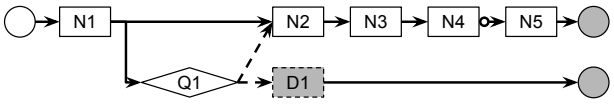

Figure 2. A D^{2}NN example. Input and output nodes are drawn as circles with the output nodes shaded. Function nodes are drawn as rectangles (regular nodes) or diamonds (control nodes). Dummy nodes have no computation and are shaded. Data edges are drawn as solid arrows and control edges as dashed arrows. A data edge with a user defined default value is decorated with a circle.

D^{2}NN Semantics Given a D^{2}NN and a particular input, we perform inference by traversing the graph starting from the input nodes. Because a D^{2}NN is a DAG, we can execute each node in a topological order (the parents of a node are ordered before it), same as conventional DNNs except a major difference: the control nodes can cause the computation of some nodes to be skipped.

After we execute a control node, it outputs a set of "control scores", one for each of its outgoing control edges. The control edge with the highest score is "activated", meaning that the node being controlled is allowed to execute. The rest of the control edges are not activated, and their controllees are not allowed to execute. For example, in Fig 2, the node Q1 controls N2 and D1. Either N2 or D1 will execute depending on which has the higher control score. A special case is that if a node has multiple incoming control edges (i.e. controlled by multiple controllers), it will execute if any of the control edges is activated.

If the execution of a node is skipped, its output will be the default value if one exists, or will be null otherwise. If the output is the default, subsequent execution will continue as usual based on the default value. If the output is null, any downstream nodes that depend on this output will in turn skip execution and have a null output unless a default value has been set. This "null" effect will propagate to the rest of the graph. For example, in Fig. 2, if N2 skips execution and outputs null, so will N3 and N4. But N5 will execute regardless because its input data edge has a default value and will never be null.

We can summarize the semantics of D^{2}NNs as follows: a D^{2}NN executes the same way as a conventional DNN except that there are control edges that can cause some nodes to be skipped. A control edge is active if and only if it has the highest score among all outgoing control edges from a node. A node is skipped if it has incoming control edges and none of them is active, or if one of its inputs is null. If a node is skipped, its output will be either null or a user-

defined default value. A null will cause downstream nodes to be skipped whereas a default value will not.

A D^{2}NN can also be thought of as a program with conditional statements. Each data edge is equivalent to a variable that is initialized to either a default value or null. Executing a function node is equivalent to executing a command assigning the output of the function to the variable. A control edge is equivalent to a boolean variable initialized to False. A control node is equivalent to a "switch-case" statement that computes a score for each of the boolean variables and sets the one with the largest score to True. Checking the conditions to determine whether to execute a function is equivalent to enclosing the function with an "if-then" statement. A conventional DNN is essentially a program with only function calls and variable assignment without any conditional statements, whereas a D^{2}NN introduces conditional statements with the conditions themselves generated by learnable functions.

4. D^{2}NN Learning

Due to the control nodes, a D^{2}NN cannot be trained the same way as a conventional DNN. In particular, the output of the network cannot be expressed as a differentiable function of all trainable parameters, especially the parameters in the control nodes. As a result, backpropagation cannot be directly applied.

The main difficulty lies in the control nodes, whose outputs are discretized into control decisions. This is similar to the situation with hard attention models in prior work [25, 3], which has used reinforcement learning. Here we adopt the same general strategy.

Learning a Single Control Node For simplicity of exposition we start with a special case where there is only one control node. We further assume that all parameters except those of this control node have been learned and fixed. That is, the goal is to learn the parameters of the control node to maximize a user-defined reward, which in our case is a combination of accuracy and efficiency. This results in a classical reinforcement learning setting: learning a control policy to take actions so as to maximize reward.

We base our learning method on Q-learning [26, 31]. We let each outgoing edge of a control node represent an action, and let the control node approximate the action-value (Q) function, which is the expected return of an action given the current state (the current input to the control node).

It is worth noting that unlike many prior works that use deep reinforcement learning, a D^{2}NN is not recurrent. In particular, for each input to the network (e.g. an image), each control node only executes once. In addition, the decisions of a control node completely depend on the current input. As a result, an action taken on one input has no effect on another input. In other words, our reinforcement learning task is "episodic" and each episode consists of only one time step. This fact simplifies our Q-learning objective to that of the following regression task:

$$L = (Q(\boldsymbol{s}, \boldsymbol{a}) - r)^2, \tag{1}$$

where r is a user-defined reward, $\boldsymbol{a}$ is an action, $\boldsymbol{s}$ is the input to control node, and Q is computed by the control node. As we can see, training a control node here is the same as training a network to predict the reward for each action under an L2 loss. In particular, we use mini-batch gradient descent; for each training example in a mini-batch, we pick the action with the largest Q, execute the rest of the network, observe a reward, and perform backpropagation using the L2 loss in Eqn. 1.

During training we also perform ϵ-greedy exploration — instead of always choosing the action with the best Q value, we choose a random action with probability ϵ. The hyperparameter ϵ is initialized to 1 and decreases over time.

The reward r is user defined. Since our goal is to optimize the trade-off between accuracy and efficiency, in our experiments we define the reward as a combination of an accuracy metric A (for example, F-score) and an efficiency metric E (for example, the inverse of the number of multiplications), that is, $\lambda A + (1 - \lambda)E$ where λ balances the trade-off.

Mini-Bags for Set-Based Metrics Our training algorithm so far has defined the state as a single training example (e.g. a single image), i.e., the control node takes actions and observe rewards on each training example independent of others. This setup, however, introduces a difficulty for optimizing for accuracy metrics that cannot be decomposed over individual examples.

Consider *precision* in the context of binary classification. Given predictions on a set of examples and the ground truth, precision is defined as the proportion of true positives among the predicted positives. Although precision can be defined on a single example, precision on a set of examples does not generally equal the average of the precisions of individual examples. In other words, precision as a metric does not decompose over individual examples and can only be computed using a set of examples *jointly*. This is different from decomposable metrics such as *error rate*, which can be computed as the average of the error rates of individual examples. In other words, if we use precision as our accuracy metric, it is not clear how to define a reward independently for each example such that maximizing this reward independently for each example would optimize the overall precision.

In general, for many accuracy metrics, including precision, recall, and F-score, we cannot compute them on individual examples and average the results. Instead, we must compute them using a set of examples as a whole. We call such metrics "set-based metrics". Our learning setup so far

is ill-equipped for such metrics because a reward is defined on each example independently.

To address this issue we generalize the definition of a state from a single input to a set of inputs. We call such a set of inputs a "mini-bag". Now, any set-based metric can be computed on a mini-bag and can be used to directly define a reward. And an action on a mini-bag $\mathbf{s} = (s_1, \ldots, s_m)$ is now a joint action $\mathbf{a} = (a_1, \ldots, a_m)$ consisting of individual actions a_i on example s_i. Let $\mathbb{Q}(\mathbf{s}, \mathbf{a})$ be the "joint" action-value function on the mini-bag $\mathbf{s}$ and the joint action $\mathbf{a}$. And we constrain the parametric form of $\mathbf{Q}$ to decompose over individual examples:

$$\mathbf{Q} = \sum_{i=1}^{m} Q(s_i, a_i), \tag{2}$$

where $Q(s_i, a_i)$ is a score given by the control node when choosing the action a_i for example s_i. We then define our new learning objective on a mini-bag of size m as

$$L = (r - \mathbf{Q}(\mathbf{s}, \mathbf{a}))^2 = (r - \sum_{i=1}^{m} Q(s_i, a_i))^2, \tag{3}$$

where r is the reward observed by choosing the joint action $\mathbf{a}$ on mini-bag $\mathbf{s}$. That is, the control node predicts an action-value for each example such that their sum approximates the reward defined on the whole mini-bag.

It is worth noting that the decomposition of $\mathbf{Q}$ into sums (Eqn. 2) enjoys a nice property: the best joint action a^* for a mini-bag under the joint action-value $\mathbf{Q}(\mathbf{s}, \mathbf{a})$ is simply the concatenation of the best actions for individual examples because maximizing the sum

$$\mathbf{a}^* = \arg\max_{\mathbf{a}}(\mathbf{Q}(\mathbf{s}, \mathbf{a})) = \arg\max_{\mathbf{a}}(\sum_{i=1}^{m} Q(s_i, a_i)) \tag{4}$$

is equivalent to maximizing the individual summands:

$$a_i^* = \arg\max_{a_i} Q(s_i, a_i), i = 1, 2...m. \tag{5}$$

That is, although we train using mini-batches of mini-bags (i.e. multiple mini-bags in a batch), during test time we still perform inference on each test example independently.

Another implication of the mini-bag formulation is that the gradient for each example is now dependent on other examples in the same mini-bag:

$$\frac{\partial \mathbf{Q}(\mathbf{s}, \mathbf{a})}{\partial x_i} = 2(r - \sum_{j=1}^{m} Q(s_j, a_j)) \frac{\partial Q(s_i, a_i)}{\partial x_i}, \tag{6}$$

where x_i is the output of any internal neuron for example i in the mini-bag. Eqn. 6 shows that there is no change to the implementation of backpropagation except that we scale the gradient using the difference between the mini-bag action-value $\mathbf{Q}$ and the mini-bag reward r.

Joint Training of All Nodes We have described how to train a single control node assuming the parameters of all other nodes are fixed. We now describe how to extend this strategy to all nodes including additional control nodes as well as regular nodes.

If a D^{2}NN has multiple control nodes, we simply train them together. For each mini-bag, we perform backpropagation for multiple losses together. Specifically, we perform inference using the current parameters, observe a reward for the whole network, and then use the same reward (which is a result of the actions of all control nodes) to backprop for each control node.

For regular nodes, we can place losses on them the same as we do on conventional DNNs. And we perform backpropagation on these losses together with the control nodes. The implementation of backpropagation is the same as conventional DNNs except that each training example may have a different network topology, that is, not all examples go through the same execution sequence. And if a node is skipped for a particular training example, then the node does not have a gradient from the example, i.e. the gradient is set to zero.

It is worth noting that our D^{2}NN framework allows arbitrary losses to be used for training regular nodes. For example, for classification we can use the cross-entropy loss on a regular node. One important detail is that the losses on regular nodes need to be properly weighted against the losses on the control nodes; otherwise the regular losses may dominate, rendering the control nodes ineffective. One way to eliminate this issue is to use Q-learning losses on regular nodes as well, i.e. treating the outputs of a regular node as action-values whenever applicable. For example, instead of using the cross-entropy loss on the ouptutted classification scores, we treat the classification scores as action-values—an estimated reward of each classification decision. This way Q-learning is applied to all nodes in a unified way and no additional hyperparameters are needed to balance different kinds of losses. In our experiments unless otherwise noted we adopt this unified approach.

5. Experiments

Implementation We implement the D^{2}NN framework in Torch [1]. Torch already provides implementations of conventional neural network modules (nodes). So a user can specify the subnetwork architecture inside a control node or a regular node using existing Torch functionalities. Our framework then handles the communication between the user-defined nodes in the forward and backward pass. For more details please see the Appendix.

High-Low Capacity D^{2}NN Our first experiment is with

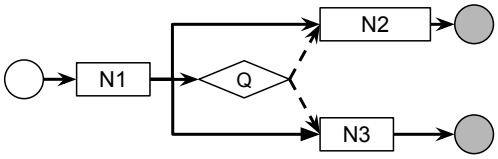

Figure 3. High-low capacity D^{2}NN. The control node Q chooses between a low-capacity subnetwork N2 and a high-capacity subnetwork N3.

a very simple D^{2}NN architecture that we call "high-low capacity D^{2}NN". It consists of a single control nodes (Q) and three regular nodes (N1,N2,N3) as illustrated in Fig.3. The control node Q makes choices between a high-capacity node N2 and a low-capacity node N3; the low-capabity node has fewer neurons and uses less computation (please see the detailed configurations of each node in the Appendix). The intuition is that the control Q node can save computation by choosing the low-capacity node for easy examples.

We test this hypothesis using a simple binary classification task in which the network classifies an input image as face or non-face. We use the Labeled Faces in the Wild [19, 20] dataset. Specifically, we use the 13k ground truth face crops (112×112 pixels) as positive examples and randomly sample 130k background crops as negative examples [1]. We hold out 11k images for validation and 22k for testing. We refer to this dataset as LFW-B and use it as a simple testbed to validate the effectiveness of our new D^{2}NN framework.

To evaluate performace we measure accuracy using the F1 score (the harmonic mean of precision and recall), a better metric than percentage of correct predictions for an unbalanced dataset. We measure computational cost using the number of actual multiplications (control nodes and regular nodes), normalized by a conventional DNN consisting of only N1 and N2, that is, the high-capacity execution path.

During training we define the Q-learning reward as a linear combination of accuracy A and efficiency E (negative cost): $r = \lambda A + (1-\lambda)E$ where $\lambda \in [0,1]$. We train instances of high-low capacity D^{2}NNs using different λ's. As λ increases, the learned D^{2}NN trades off efficiency for accuracy. Fig. 6(a) plots the accuracy-cost curve on the test set; it also plots the accuracy and efficiency achieved by a conventional DNN with only the high capacity path N1+N2 (High NN) and a conventional DNN with only the low capacity path N1+N3 (Low NN).

As we can see, the D^{2}NN achieves a trade-off curve close to the upperbound: there are points on the curve that are as fast as the low-capacity node and as accurate as the high-capacity node. Fig. 5(left) plots the distribution of examples going through different execution paths. It shows that as λ increases, accuracy becomes more important and more examples go through the high-capacity node. These results suggest that our D^{2}NN learning algorithm is effective for networks with a single control node.

[1] background crops are those that have an intersection over union less than 0.3.

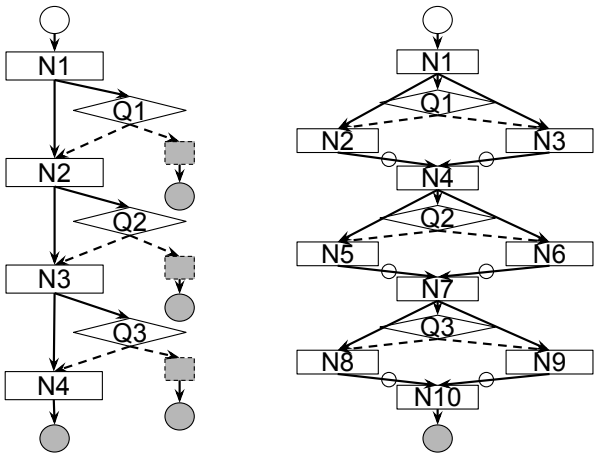

Figure 4. Left: a cascade D^{2}NN. Right: a chain D^{2}NN.

Cascade D^{2}NN We next experiment with a more sophisticated design that we call a "cascade D^{2}NN" (Fig. 4 Left). The cascade D^{2}NN consists of a sequence of four regular nodes (N1 to N4) and three control nodes (Q1-Q3) [2]. Each control node decides whether to continue executing the next regular node or to halt (i.e. executing a dummy node). This architecture is inspired by the standard cascade design commonly used in computer vision. The intuition is that many negative examples may be rejected early using simple features.

We evaluate the network on the same LFW-B face classification task using the same evaluation protocol as we did for the high-low capacity D^{2}NN. Fig. 6(b) plots the accuracy-cost tradeoff curve for the full model (4N3Q) as well as the curves for ablated models with fewer control nodes (4N2Q and 4N1Q). Also included are the accuracy and cost of a conventional DNN (4N) consisting of the same regular nodes. We can see that the cascade D^{2}NN can achieve a close to optimal trade-off, reducing computation with negligible loss of accuracy. In addition, we see that adding more control nodes provides more options for early rejection and achieves significantly better curves. This result demonstrates that our algorithm is successful for jointly training multiple control nodes.

Chain D^{2}NN Our third design is a "Chain D^{2}NN" (Fig. 4 Right) [3]. The network is shaped as a chain, where each link consists of a control node selecting between two regular nodes. In other words, we perform a sequence of

[2] Please see configuration details in the Appendix.

[3] Please see configuration details in the Appendix.

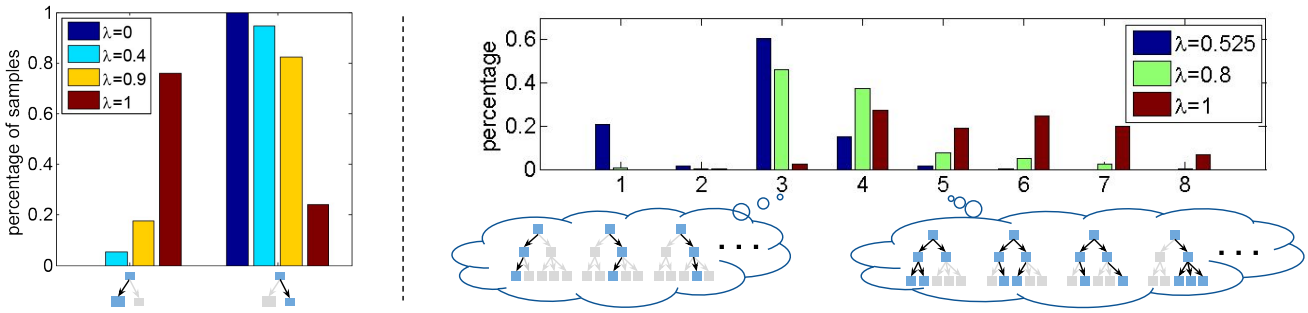

Figure 5. Distribution of examples going through different execution paths. Skipped nodes in a D^{2}NN are in grey. The hyperparameter λ controls the trade-off between accuracy and efficiency. A bigger λ gives more weight to accuracy. *Left*: distributions for the high-low capacity D^{2}NN. *Right*: distributions for the hierarchical D^{2}NN. The X-axis is the number of nodes activated.

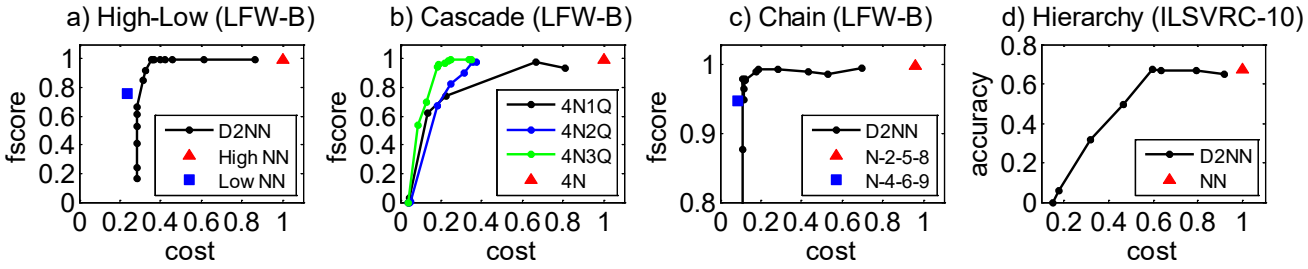

Figure 6. The accuracy-cost curves of various D^{2}NN architectures, as well as conventional DNN baselines consisting of only regular nodes.

vector-to-vector transforms; for each transform we choose between two subnetworks possibly with different capacities. This chain D^{2}NN is qualitatively different from a cascade or other D^{2}NNs with a tree-shaped *data graph* because the chain D^{2}NN allows two divergent *data paths* to merge again. That is, the number of possible execution paths can be exponential in the number of D^{2}NN nodes—for a chain of length n, there are 2^n possible execution paths.

In Fig. 4(right), the first link has a low-capacity node (N2) and a high-capacity node (N3). If one of them is chosen, the other will output a default value zero. The node N4 simply adds the outputs of N2 and N3 together. Fig. 6(c) plots the accuracy-cost curve on the LFW-B face classification task. In the same figure we also compare with two baselines: a conventional DNN consisting of only the low capacity nodes (N2, N5, N8), and a conventional DNN consisting of only the high capacity nodes (N3, N6, N9). Both conventional DNNs are trained with the cross-entropy loss. The cost is measured as the number of multiplications, normalized by the cost of the high-capacity baseline (N-3-6-9).

Fig. 4(c) shows that the chain D^{2}NN achieves a trade-off curve close to the optimal and can speed up computatoin significantly with very little accuracy loss. This shows that our learning algorithm is effective for a D^{2}NN whose data graph is a general directed acyclic graph instead of a tree.

Hierarchical D^{2}NN In this experiment we design a D^{2}NN for hierarchical multiclass classification. The idea is to first classify images to coarse categories and then to fine categories. This idea has been explored by numerous prior works [24, 5, 9], but here we show that the same idea can be implemented via a D^{2}NN trained end to end.

We use ILSVRC-10, a subset of the ILSVRC-65 dataset [8]. In our ILSVRC-10, there are 10 classes organized into a 3-layer hierarchy: 2 superclasses, 5 coarse classes and 10 leaf classes [4]. Each class has 500 training images, 50 validation images, and 150 test images.

Fig. 7(right) illustrates the design of our hierarchical D^{2}NN [5]. The hierarchy in this D^{2}NN mirrors the semantic hierarchy in ILSVRC-10. An image first goes through the root node N1. Then the control node Q1 decides whether to descend the left branch of the hierarchy (N2 and its children), and the control node Q2 decides whether to descend the right branch of the hierarchy (N3 and its children). The leaf nodes N4 to N8 are each responsible for classifying two fine-grained leaf classes.

It is important to note that an input image can go down parallel paths in the hierarchy. That is, it can descend both the left branch and the right branch because Q1 and Q2 make separately decisions. This "multi-threading" allows the network to avoid committing to a single path prematurely if an input image is difficult or ambiguous.

Fig. 6(d) plots the accuracy-cost curve of our hierarchical D^{2}NN. The accuracy is measured as the proportion of correctly classified test examples. The cost is measured as

[4] See the Appendix for more details

[5] Please see configuration details in the Appendix.

the number of multiplications, normalized by the cost of a conventional DNN consisting only of the regular nodes (denoted as NN in the figure). We can see in Fig. 6(d) that the hierarchical D^{2}NN can match the accuracy of the full network with about half of the computational cost.

Fig. 5(right) plots for the hierarchical D^{2}NN the distribution of examples going through execution sequences with different numbers of nodes activated. Due to the parallelism of D^{2}NN, there can be many different execution sequences. We also see that as λ increases, accuracy is given more weight and more nodes are activated.

Comparison with Dynamic Capacity Networks In this experiment we empirically compare our approach to closely related prior work. Note that this comparison is not always practical because many prior works either do not report efficiency [11] or measure efficiency using real time [23, 4], which depends on hardware and is hard to reproduce.

Here we compare D^{2}NNs with Dynamic Capacity Networks (DCN) [2], for which efficency measurements are available as the number of multiplications. Given an input image, a DCN applies an additional high capacity subnetwork to a set of image patches, selected using a hand-designed saliency based policy. The idea is that more intensive processing is only necessary for certain image regions.

To compare, we evaluate on the same multiclass classification task. We use the same dataset Cluttered MNIST [25], which consists of MNIST digits randomly placed on a background cluttered with fragments of other digits. We train a chain D^{2}NN of length 4 [6] as shown in Fig. 7(left), which implements the same idea of choosing a high-capacity alternative subnetwork for certain inputs. Fig. 8 plots the accuracy-cost curve of our D^{2}NN as well as the accuracy-cost point achieved by the DCN in [2]—an accuracy of 0.9861 and and a cost of 2.77×10^7. The closet point on our curve is an slightly lower accuracy of 0.9698 but slightly better efficiency (a cost of 2.66×10^7). Note that although our accuracy of 0.9698 is lower, it compares favorably to those of other state of the art methods such as DRAW (0.9189) [15] and RAM (0.9664) [25].

Effect of Mini-Bag Size Our last experiment concerns a new hyperparameter specific to our learning algorithm: the size of a mini-bag. Recall that mini-bags are introduced to optimize set-based metrics such as the F1 score. In this experiment we empirically validate this algorithmic design. We train the same high-low capacity D^{2}NN on the LFW-B task but vary the size of the mini-bags (1,4,8,16,32,128) while fixing the total number of examples in a mini-batch to 128. That is, a mini-bag size of 8 leads to 16 mini-bags in a mini-batch. A mini-bag size of 1 is the same as conventional mini-batch training. Fig. 9 plots the accuracy-cost curves achieved by different mini-bag sizes, showing the best mini-bag size is 8, i.e. using multiple mini-bags in a mini-batch indeed improves results.

[6] Please see configuration details in the Appendix.

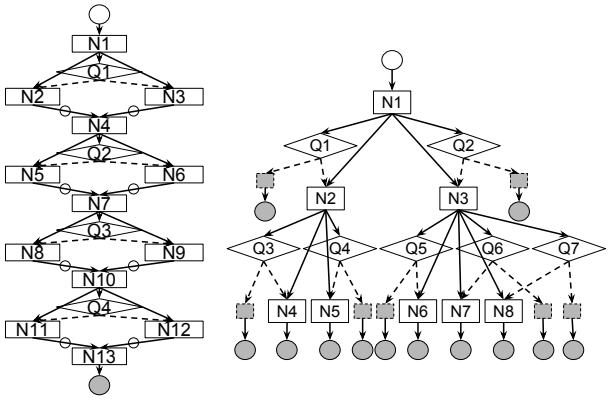

Figure 7. Left: A chain D^{2}NN for the Cluttered MNIST dataset. Right: A Hierarchical D^{2}NN.

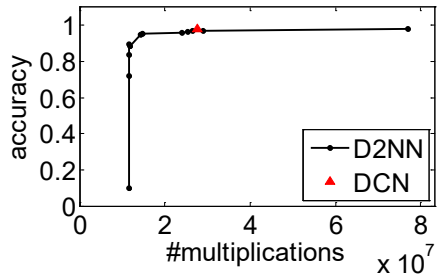

Figure 8. Accuracy-cost curve for a chain D^{2}NN on the CMNIST task compared to Dynamic Capacity networks (DCN) [2].

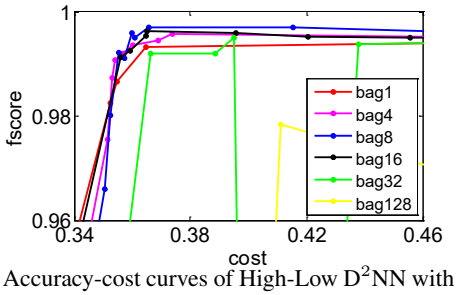

Figure 9. Accuracy-cost curves of High-Low D^{2}NN with different mini-bag sizes.

6. Conclusion

We have introduced Dynamic Deep Neural Networks (D^{2}NN), a new type of feed-forward deep neural networks that allow selective execution. Extensive experiments have demonstrated that D^{2}NNs are flexible and effective for optimizing accuracy-efficiency trade-offs.

7. Appendix

7.1. Implementation Details

We implement the D^{2}NN framework in Torch [1]. Torch already provides implementations of conventional neural network modules (nodes). So a user can specify the sub-network architecture inside a control node or a regular node using existing Torch functionalities. Our framework then handles the communication between the user-defined nodes in the forward and backward pass.

To handle parallel paths, default-valued nodes and nodes with multiple data parents, we need to keep track of a example's execution status (which nodes are activated by this example) and output status (which nodes have output for this example). A example's output status is different from its execution status if some nodes are not activated but have default values. For runtime efficiency, we implement the tracking of examples at the mini-batch level. That is, we perform forward and backward passes for a mini-batch of examples as a regular DNN does. Each mini-batch consists of several mini-bags of images.

We describe the implementation of D^{2}NN learning procedure as two steps. First, the preprocessing step: When a user-defined D^{2}NN model is fed into our framework, we first perform a breadth-first search to get the DAG orders of nodes while performing structure error checks, contructing data and control relationships between nodes and calculating the cost (number of multiplications) of each node.

After the preprocessing, the training step is similar to a regular DNN: a forward pass and a backward pass. All nodes are visited according to a topological ordering in a forward pass and the reverse ordering in a backward pass.

For each function node, the forward pass has three steps: fetch inputs, forward inside the node, and send data or control signals to children nodes. When dealing with multiple data inputs and multiple control signals, the D^{2}NN will filter examples with more than one null inputs or all negative control signals. When a default value has been set for a node, all examples have to send out data. If the node is not activated for a particular example, the output will take the default value. A backward pass has similar logic: fetch gradients from children, perform the backward pass inside and send out gradients to parents. It is worth noting that when a default value is used in a node, the gradients can be blocked by this node because it is not actually executed.

7.2. ILSVRC-10 Semantic Hierarchy

The ILSVRC-10 dataset is a subset of the ILSVRC-65 dataset [8]. In our ILSVRC-10, there are 10 classes organized into a 3-layer hierarchy: 2 superclasses, 5 coarse classes and 10 leaf classes as in Fig 10. Each class has 500 training images, 50 validation images, and 150 test images.

7.3. Configurations

High-Low Capacity D^{2}NN The high-low capacity D^{2}NN consists of a single control node (Q) and three regular nodes (N1,N2,N3) as illustrated in Fig. 3.

- Node N1: a convolutional layer with a 3×3 filter size, 8 filters and a stride of 2, followed by a 3×3 max-pooling layer with a stride of 2.
- Node N2: a convolutional layer with a 3×3 filter size and 16 filters, followed by a 3×3 max-pooling layer with a stride of 2. The output is reshaped and fed into a fully connected layer with 512 neurons followed by another fully connected layer with the 2-class output.
- Node N3: three 3×3 max-pooling layers, each with a stride of 2, followed by two fully connected layers with 32 neurons and the 2-class output.
- Node Q1: a convolutional layer with a 3×3 filter size and 2 filters, followed by a 3×3 max-pooling layer with a stride of 2. The output is reshaped and fed into a fully connected layer with 128 neurons followed by another fully connected layer with the 2-action output.

Cascade D^{2}NN The cascade D^{2}NN consists of a sequence of four regular nodes (N1 to N4) and three control nodes (Q1-Q3) as in Fig. 4(left).

- Node N1: a convolutional layer with a 3×3 filter size, 2 filters and a stride of 2, followed by a 3×3 max-pooling layer with a stride of 2.
- Node N2: a convolutional layer with a 3×3 filter size and 8 filters, followed by a 3×3 max-pooling layer with a stride of 2.
- Node N3: a convolutional layer with a 3×3 filter size and 32 filters.
- Node N4: two convolutional layers with both 3×3 filter sizes and 32, 64 filters repectively, augmented by a 3×3 max-pooling layer with a stride of 2. The output is reshaped and fed into a fully connected layer with 512 neurons followed by another fully connected layer with the 2-class output.
- Node Q1: three 3×3 max-pooling layers with strides of 2. The output is reshaped and fed into a fully connected layer with the 2-action output.
- Node Q2: two 3×3 max-pooling layers with strides of 2, followed by a convolutional layer with a 1×1 filter size and 2 filters. The output is reshaped and fed into a fully connected layer with the 2-action output.

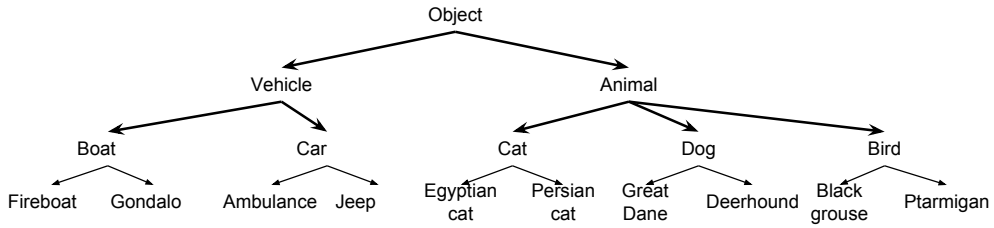

Figure 10. The semantic class hierarchy of the ILSVRC-10 dataset.

- Node Q3: a 3×3 max-pooling layer with a stride of 2. The output is reshaped and fed into a fully connected layer with the 2-action output.

Chain D^{2}NN The Chain D^{2}NN is shaped as a chain, where each link consists of a control node selecting between two regular nodes. In the experiments of LFW-B dataset, we use a 3-stage Chain D^{2}NN as in Fig. 4(right).

- Node N1: a convolutional layer with a 3×3 filter size, 2 filters and a stride of 2, followed by a 3×3 max-pooling layer with a stride of 2.
- Node N2: a convolutional layer with a 1×1 filter size and 16 filters.
- Node N3: a convolutional layer with a 3×3 filter size and 16 filters.
- Node N4: a 3×3 max-pooling layer with a stride of 2.
- Node N5: a convolutional layer with a 1×1 filter size and 32 filters.
- Node N6: two convolutional layers with both 3×3 filter sizes and 32, 32 filters repectively.
- Node N7: a 3×3 max-pooling layer with a stride of 2.
- Node N8: a convolutional layer with a 1×1 filter size and 32 filters followed by a 3×3 max-pooling layer with a stride of 2. The output is reshaped and fed into a fully connected layer with 256 neurons.
- Node N9: a convolutional layer with a 3×3 filter size and 64 filters. The output is reshaped and fed into a fully connected layer with 256 neurons.
- Node N10: a fully connected layer with the 2-class output.
- Node Q1: a convolutional layer with a 3×3 filter size and 8 filters with a 3×3 max-pooling layer with a stride of 2 before and a 3×3 max-pooling layer with a stride of 2 after. The output is reshaped and fed into two fully connected layers with 64 neurons and the 2-action output respectively.
- Node Q2: a 3×3 max-pooling layer with a stride of 2 followed by a convolutional layer with a 3×3 filter size and 4 filters. The output is reshaped and fed into two fully connected layers with 64 neurons and the 2-action output respectively.
- Node Q3: a convolutional layer with a 3×3 filter size and 2 filters. The output is reshaped and fed into two fully connected layers with 64 neurons and the 2-action output respectively.

Hierarchical D^{2}NN Fig. 7(right) illustrates the design of our hierarchical D^{2}NN.

- Node N1: a convolutional layer with a 11×11 filter size, 64 filters, a stride of 4 and a 2×2 padding, followed by a 3×3 max-pooling layer with a stride of 2.
- Node N2 and N3: a convolutional layer with a 5×5 filter size, 96 filters and a 2×2 padding.
- Node N4 N8: a 3×3 max-pooling layer with a stride of 2 followed by three convolutional layers with 3×3 filter sizes and 160, 128, 128 filters respectively. The output is fed into a 3×3 max-pooling layer with a stride of 2 and three fully connected layers with 2048 neurons, 2048 neurons and the 2 fine-class output respectively.
- Node Q1 and Q2: two convolutional layers with 5×5, 3×3 filter sizes and 16, 32 filters respectively (the former has a 2×2 padding), each followed by a 3×3 max-pooling layer with a stride of 2. The output is reshaped and fed into three fully connected layers with 1024 neurons, 1024 neurons and the 2-action output respectively.
- Node Q3 Q7: two convolutional layers with 5×5, 3×3 filter sizes and 16, 32 filters respectively (the former has a 2×2 padding), each followed by a 3×3 max-pooling layer with a stride of 2. The output is reshaped and fed into three fully connected layers with 1024 neurons, 1024 neurons and the 2-action output respectively.

Comparison with Dynamic Capacity Networks We train a chain D^{2}NN of length 4 as in Fig. 7 (left).

- Node N1: a convolutional layer with a 3×3 filter size and 24 filters.
- Node N3: a convolutional layer with a 3×3 filter size and 24 filters.
- Node N4: a 2×2 max-pooling layer with a stride of 2.
- Node N6: a convolutional layer with a 3×3 filter size and 24 filters.
- Node N7: an identity layer which directly uses inputs as outputs.
- Node N9: a convolutional layer with a 3×3 filter size and 24 filters.
- Node N10: a 2×2 max-pooling layer with a stride of 2.
- Node N12: a convolutional layer with a 3×3 filter size and 24 filters.
- Node N2, N5, N8, N11: an identity layer.
- Node N13: a convolutional layer with a 4×4 filter size, 96 filters, a stride of 2 and no padding, followed by a 11×11 max-pooling layer. The output is reshaped and fed into a fully connected layer with the 10-class output.
- Node Q1: a convolutional layer with a 3×3 filter size and 8 filters with two 2×2 max-pooling layers with strides of 2 before and one 2×2 max-pooling layer with a stride of 2 after. The output is reshaped and fed into two fully connected layers with 256 neurons and the 2-action output respectively.
- Node Q2: a convolutional layer with a 3×3 filter size and 8 filters with a 2×2 max-pooling layer with a stride of 2 before and a 2×2 max-pooling layer with a stride of 2 after. The output is reshaped and fed into two fully connected layers with 256 neurons and the 2-action output respectively.
- Node Q3: a convolutional layer with a 3×3 filter size and 8 filters with a 2×2 max-pooling layer with a stride of 2 before and a 2×2 max-pooling layer with a stride of 2 after. The output is reshaped and fed into two fully connected layers with 256 neurons and the 2-action output respectively.
- Node Q4: a convolutional layer with a 3×3 filter size and 8 filters, followed by a 2×2 max-pooling layer with a stride of 2. The output is reshaped and fed into two fully connected layers with 256 neurons and the 2-action output respectively.

For all 5 D^{2}NNs, all convolutional layers use 1×1 padding and each is followed by a ReLU layer unless specified individually. Each fully connected layer except the output layers is followed by a ReLU layer.

References

[1] Torch. http://torch.ch/. 5, 9

[2] A. Almahairi, N. Ballas, T. Cooijmans, Y. Zheng, H. Larochelle, and A. C. Courville. Dynamic capacity networks. In *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 2549–2558, 2016. 2, 8

[3] J. Ba, V. Mnih, and K. Kavukcuoglu. Multiple object recognition with visual attention. *arXiv preprint arXiv:1412.7755*, 2014. 2, 4

[4] E. Bengio, P.-L. Bacon, J. Pineau, and D. Precup. Conditional computation in neural networks for faster models. *arXiv preprint arXiv:1511.06297*, 2015. 2, 8

[5] S. Bengio, J. Weston, and D. Grangier. Label embedding trees for large multi-class tasks. In *Advances in Neural Information Processing Systems*, pages 163–171, 2010. 2, 7

[6] Y. Bengio. Deep learning of representations: Looking forward. In *International Conference on Statistical Language and Speech Processing*, pages 1–37. Springer, 2013. 2

[7] Y. Chen, T. Luo, S. Liu, S. Zhang, L. He, J. Wang, L. Li, T. Chen, Z. Xu, N. Sun, et al. Dadiannao: A machine-learning supercomputer. In *Microarchitecture (MICRO), 2014 47th Annual IEEE/ACM International Symposium on*, pages 609–622. IEEE, 2014. 2

[8] J. Deng, J. Krause, A. C. Berg, and L. Fei-Fei. Hedging your bets: Optimizing accuracy-specificity trade-offs in large scale visual recognition. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3450–3457. IEEE, 2012. 7, 9

[9] J. Deng, S. Satheesh, A. C. Berg, and F. Li. Fast and balanced: Efficient label tree learning for large scale object recognition. In *Advances in Neural Information Processing Systems*, pages 567–575, 2011. 2, 7

[10] M. Denil, L. Bazzani, H. Larochelle, and N. de Freitas. Learning where to attend with deep architectures for image tracking. *Neural computation*, 24(8):2151–2184, 2012. 2

[11] L. Denoyer and P. Gallinari. Deep sequential neural network. *arXiv preprint arXiv:1410.0510*, 2014. 2, 8

[12] E. L. Denton, W. Zaremba, J. Bruna, Y. LeCun, and R. Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *Advances in Neural Information Processing Systems*, pages 1269–1277, 2014. 2

[13] P. F. Felzenszwalb, R. B. Girshick, and D. McAllester. Cascade object detection with deformable part models. In *Computer vision and pattern recognition (CVPR), 2010 IEEE conference on*, pages 2241–2248. IEEE, 2010. 2

[14] V. Gokhale, J. Jin, A. Dundar, B. Martini, and E. Culurciello. A 240 g-ops/s mobile coprocessor for deep neural networks. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2014 IEEE Conference on*, pages 696–701. IEEE, 2014. 2

[15] K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra. Draw: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015. 2, 8

[16] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan. Deep learning with limited numerical precision. *arXiv preprint arXiv:1502.02551*, 2015. 2

[17] S. Han, J. Pool, J. Tran, and W. Dally. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, pages 1135–1143, 2015. 2

[18] B. Hassibi and D. G. Stork. *Second order derivatives for network pruning: Optimal brain surgeon*. Morgan Kaufmann, 1993. 2

[19] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007. 6

[20] G. B. H. E. Learned-Miller. Labeled faces in the wild: Updates and new reporting procedures. Technical Report UM-CS-2014-003, University of Massachusetts, Amherst, May 2014. 6

[21] V. Lebedev, Y. Ganin, M. Rakhuba, I. Oseledets, and V. Lempitsky. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. *arXiv preprint arXiv:1412.6553*, 2014. 2

[22] Y. LeCun, J. S. Denker, S. A. Solla, R. E. Howard, and L. D. Jackel. Optimal brain damage. In *NIPs*, volume 89, 1989. 2

[23] H. Li, Z. Lin, X. Shen, J. Brandt, and G. Hua. A convolutional neural network cascade for face detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5325–5334, 2015. 2, 8

[24] B. Liu, F. Sadeghi, M. Tappen, O. Shamir, and C. Liu. Probabilistic label trees for efficient large scale image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 843–850, 2013. 2, 7

[25] V. Mnih, N. Heess, A. Graves, et al. Recurrent models of visual attention. In *Advances in Neural Information Processing Systems*, pages 2204–2212, 2014. 2, 4, 8

[26] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013. 4

[27] M. J. Saberian and N. Vasconcelos. Boosting algorithms for detector cascade learning. *Journal of Machine Learning Research*, 15(1):2569–2605, 2014. 2

[28] M. F. Stollenga, J. Masci, F. Gomez, and J. Schmidhuber. Deep networks with internal selective attention through feedback connections. In *Advances in Neural Information Processing Systems*, pages 3545–3553, 2014. 2

[29] C. Sun, M. Paluri, R. Collobert, R. Nevatia, and L. Bourdev. Pronet: Learning to propose object-specific boxes for cascaded neural networks. *arXiv preprint arXiv:1511.03776*, 2015. 2

[30] Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for facial point detection. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 3476–3483. IEEE, 2013. 2

[31] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*, volume 1. 4

[32] P. Viola and M. J. Jones. Robust real-time face detection. *International journal of computer vision*, 57(2):137–154, 2004. 2

[33] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. *arXiv preprint arXiv:1502.03044*, 2(3):5, 2015. 2

[34] X. Zhang, J. Zou, K. He, and J. Sun. Accelerating very deep convolutional networks for classification and detection. *arXiv preprint arXiv:1505.06798*, 2015. 2